An algorithm for the fitting of planet models to *Kepler* light curves

Peter Tenenbaum*[a], Stephen T. Bryson[a], Hema Chandrasekaran[a,b], Jie Li[a], Elisa Quintana[a], Joseph D. Twicken[a], Jon M. Jenkins[a]

[a]SETI Institute/NASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94035
[b]Lawrence Livermore National Laboratory, PO Box 808, L-478, Livermore, CA USA 94551

ABSTRACT

We describe an algorithm which fits model planetary system parameters to light curves from *Kepler Mission* target stars. The algorithm begins by producing an initial model of the system which is used to seed the fit, with particular emphasis on obtaining good transit timing parameters. An attempt is then made to determine whether the observed transits are more likely due to a planet or an eclipsing binary. In the event that the transits are consistent with a transiting planet, an iterative fitting process is initiated: a wavelet-based whitening filter is used to eliminate stellar variations on timescales long compared to a transit; a robust nonlinear fitter operating on the whitened light curve produces a new model of the system; and the procedure iterates until convergence upon a self-consistent whitening filter and planet model. The fitted transits are removed from the light curve and a search for additional planet candidates is performed upon the residual light curve. The fitted models are used in additional tests which identify false positive planet detections: multiple planet candidates with near-identical fitted periods are far more likely to be an eclipsing binary, for example, while target stars in which the model light curve is correlated with the star centroid position may indicate a background eclipsing binary, and subtraction of all model planet candidates yields a light curve of pure noise and stellar variability, which can be used to study the probability that the planet candidates result from statistical fluctuations in the data.

Keywords: Kepler mission, exoplanet, transit, model fit, transit photometry

1. INTRODUCTION

The *Kepler Mission*[1,2] uses a space-based photometer with a 115 square degree field of view to search for transit signatures of exoplanets, with a particular emphasis on Earth-size exoplanets in the habitable zones (HZs) of their parent stars. In order to detect the transits of Earth-sized exoplanets, *Kepler* records the flux from over 150,000 stars in its field of view at 30 minute intervals. Since the HZ of a Sun-like star corresponds to an orbital period on the order of one year, the *Kepler Mission* will observe the selected ensemble of target stars for at least 3.5 years, such that exoplanets in the HZ will generally have at least three transits in the dataset; this is necessary, since three transits are required to demonstrate that the transit occurrences are truly periodic and therefore are consistent with a transiting exoplanet.

The first step in detecting planet candidates in the *Kepler* dataset is the analysis performed by the Transiting Planet Search (TPS) module,[3] which examines the flux time series of each star and identifies periodic short-duration dips in a flux time series which are consistent with a transiting exoplanet. In a recent TPS analysis of 90 days' worth of *Kepler* data, the number of flux time series which were flagged as containing potential transiting planet candidates was a few percent of the total number of target stars. This represents a dramatic reduction in the total number of targets which must be further examined, but still leaves several thousand flux time series which contain transit-like features; those features can be as small as 100 parts per million (PPM), and in many cases are due to artifacts, anomalies, or astrophysical signatures which mimic the features of a transiting planet. The next step in *Kepler* data science processing is the Data Validation (DV) module.[4] The DV module performs

*peter.tenenbaum@nasa.gov
©2010 Society of Photo-Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

a number of tests which can be used to aid in determining whether a detected transit signature is more likely to be a transiting planet or more likely to be one class or another of false positive.

The DV module tests require a model of the planetary system which is consistent with the observed transit periods, durations, and depths. Such models are produced by a planet-model fitter which is incorporated into DV, and described in this paper.

2. PRE-FITTING ANALYSIS

The starting point for the DV fitter is the output from the TPS module. For each target star in the data set, TPS produces a Threshold Crossing Event data structure (TCE). Each TCE specifies the period, center time of first transit (also known as the "transit epoch"), and approximate duration of the set of features in the flux time series which most strongly resemble a transiting planet. While the depth of the transit is not included directly in the TCE, the strength of the transit signal is included, in units of multiples of the noise limit for the detection of a signal with the specified period and duration. This allows the DV fitter to produce an initial estimate of the transit model which seeds the fit, as described in Section 3.1.

Prior to generation of the initial transit model estimate, a number of pre-fitting steps are taken which improve the accuracy of the parameters provided in the initial TCE. These steps also allow some poorly formed transiting planet candidates to be identified and rejected prior to fitting.

2.1 Transiting Planet Search

The TPS metric for the strength of a transit signature is Multiple Event Statistic (MES): this is the ratio of the detected signature's strength to the noise limit for a transit signature of the selected period and duration, or equivalently the signal-to-noise ratio (SNR) for the detection of the series of transits.[3] For a given light curve, TPS analyzes a set of user-specified transit durations, and reports TCEs for the period and epoch which result in the maximum MES for each selected transit duration. All light curves which have a TCE with MES over a specified threshold are then passed to DV for additional analysis. In the current configuration, TPS searches for 3-hour, 6-hour, and 12-hour transit durations, and TCEs with a MES of 7.1 σ or greater are analyzed in DV.

The nature of the TPS algorithm is such that it does not attempt to ensure that all transits associated with a given MES are of comparable depth; as a result, a light curve with one or two very deep transit-like features is likely to be flagged as having a MES which is larger than the threshold. In Fig. 1, we see examples of two classes of TCEs. In the top plot is the flux time series for a true transiting planet candidate, with transit-like features which are of comparable depth. In the bottom plot is a flux time series which contains a single very large feature and a bump which is slightly above the noise floor; the TPS algorithm identifies this flux time series as containing a transit signature with a MES of 14.7 σ, despite the fact that it is actually just the artifact which contains all the meaningful contributions to the MES.

The DV fitter performs an additional screening on the TCE from TPS to identify cases which resemble the bottom plot in Fig. 1. This is accomplished by examining the Single Event Statistics (SES) which contribute to the MES, where each transit's SES is the SNR for detection of that transit (similar to the way that the MES is the SNR for detection of the series of transits). DV computes the ratio between the MES and the largest SES which contributes to the MES; if this ratio is smaller than a specified ratio, the event is considered too poor a candidate to fit, and DV moves on to its next target star.

In the limit of Gaussian noise and an infinite number of samples per single transit-like event, the relationship between the MES and the SES for a true transit signature is MES $= \sqrt{N_{\mathrm{event}}}$ SES. Since there are finite samples per transit and the noise in the light curves is non-Gaussian and non-stationary, the SES in a given MES have a distribution of values, and the largest SES value can be significantly larger than the typical value; however, it is necessary to consider the largest SES value rather than, for example, the median SES value in order to identify cases such as the one shown in the lower half of Figure 1. Given all these issues, the current cutoff for further analysis is a MES value which is at least 1.25 times the largest SES value. With the cutoff thus configured, about 50% of all TCEs which are passed to DV from TPS are rejected. The ensemble of rejected targets is overwhelmingly dominated by flux time series which contain clearly-visible artifacts which are driving the detection process, and contains relatively few flux time series which appear to the eye to contain potential

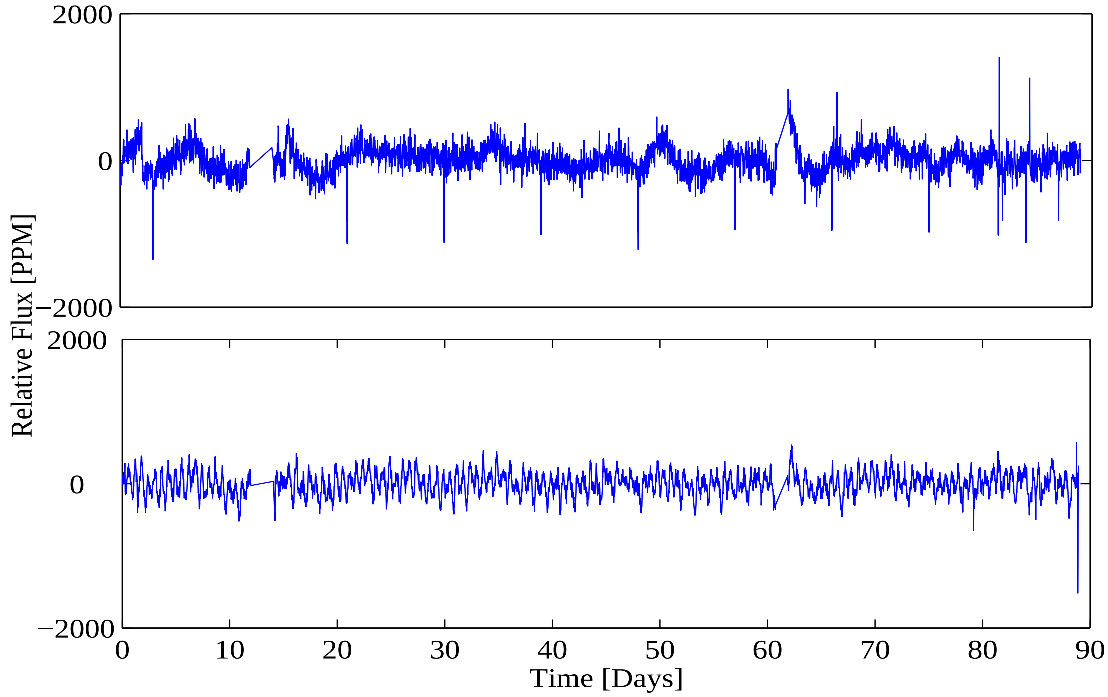

Figure 1. Two examples of light curves which produce threshold-crossing events. Top: a planet candidate, with multiple transit-like features of comparable depth. Bottom: a spurious candidate, with a single large feature at the end of the light curve and a small bump at 79 days which are combined and interpreted as a threshold crossing event by TPS.

transiting planet signatures; in short, the applied MES/SES cut mainly eliminates uninteresting targets while also eliminating a tolerably small number of interesting targets.

2.2 Transit Timing Estimate Improvement

A problem similar to the spurious TCE in Section 2.1 can cause the estimated period from TPS to be a harmonic of the actual period. Consider a flux time series as shown in the top portion of Fig. 2: a fluctuation which occurs close to the midpoint between two actual transits can cause the MES of the TCE with half the actual period to be slightly larger than the MES of the TCE which has the correct period. In this case, the TCE with half the correct period will be sent to DV, resulting in an initial estimate of the period which is far from correct. This in turn yields an extremely poor seed from which the fitting process is unlikely to be able to recover.

Upon deeper examination, the incorrect period selection is almost always caused by an unusually-shaped distribution of MES versus orbital period: the distribution shows a peak at both the incorrect and the correct period, but while the *central* value of the peak at the correct period is higher than the *central* value of the peak at the incorrect one, the peak associated with the incorrect period will have an off-center "bump" which has an even higher value; it is this "bump" which is detected by TPS and reported as the MES and orbital period for the system. A TCE with the correct period is obtained by performing an additional set of searches in TPS. Instead of performing a single search across the full range of possible periods, each additional search covers a much-reduced range of periods; the searches cover the period in the original TCE, plus the first few subharmonics of that period. In each of these searches, the MES reported is the value at the center of the detected peak, rather than the absolute maximum value; this allows the refined search to ignore an outlier MES and correctly determine the orbital period.[†] At this time, the first four subharmonics are searched, each across a window of ± 2 days. Thus, an initial TCE with a period of 10 days would trigger additional searches which examine periods of 8-12 days,

[†]Since the original development of TPS and DV, TPS has been revised to always use the peak-center value rather

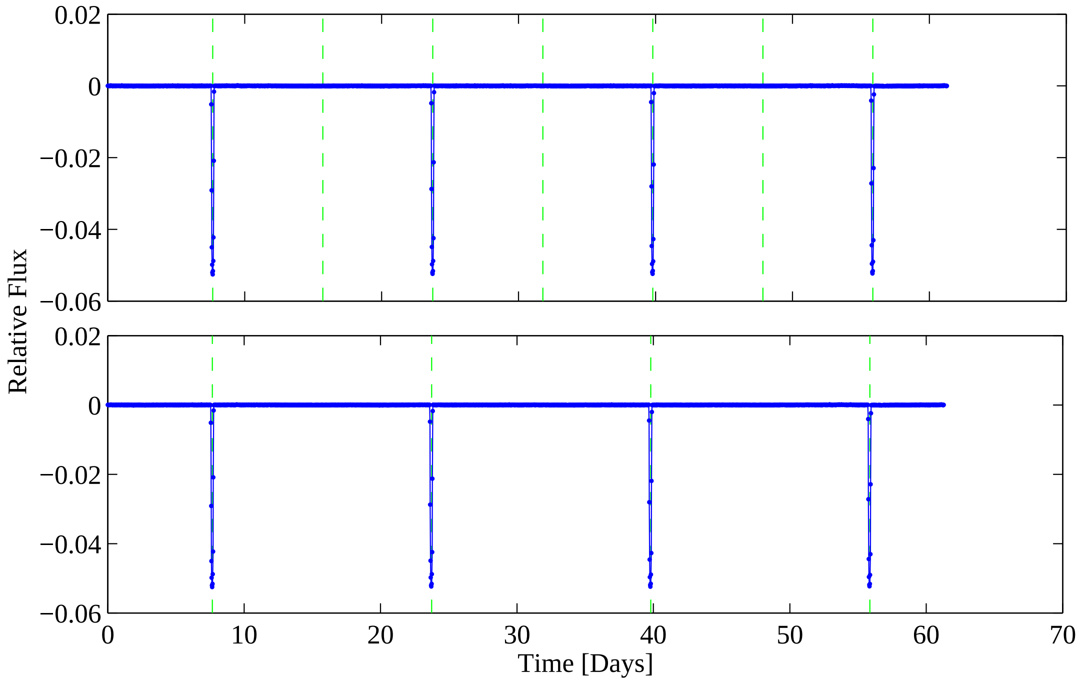

Figure 2. Example of a flux time series which requires improvement of its timing estimate. Top: flux time series with incorrect period estimate. Bottom: the same flux time series after timing estimate improvement. In both cases, the predicted transit timings are indicated by vertical green dashed lines.

18-22 days, 28-32 days, and 38-42 days. This yields a total of four new TCEs, one from each subharmonic. The new TCE with the maximum MES is then accepted as the correct one and used to seed the fit. The bottom half of Figure 2 shows the results of such a search: in this case, the second subharmonic has been correctly identified as the optimal one for the search.

In the event of an eclipsing binary in which the intervals between eclipses are close to rational (for example, 2:1 or 3:1), the procedure outlined above will still not converge upon the correct period. For example, in a 2:1 case, the third harmonic of the correct period will have the maximum MES because it overlaps all of the primary eclipses of the system, plus all of the secondary eclipses of the system, and also a set of times in which there is no eclipse at all; the correct period will overlap only the primary eclipses. In order to address this issue, the TCEs from the subharmonic searches are examined for evidence of extremely large transits at the locations predicted by the TCEs. For example, in the case of an eclipsing binary with 2:1 intervals, the transits predicted by the first and second subharmonic TCEs do not all line up with large transit events, while the transits predicted by the third subharmonic TCE will all line up with large transit events. This difference allows the correct period to be deduced even in such cases.

2.3 Eclipsing Binary Removal

In general, the DV fitter is capable of fitting a light curve in which the transit-like features are due to an eclipsing binary. In the specific case of eclipsing binaries with extremely deep transits, the fitter is generally unable to converge correctly. This issue is addressed by logic which prevents the fitter from operating in any case in which all of the transits predicted by the TCE are deeper than a threshold, which is currently set to 15%. In the event of such an eclipsing binary system, the transit parameters (epoch, period, and depth) are recorded for later use

than the absolute maximum value of MES. This has dramatically improved the ability of TPS to return a correct orbital period without any additional search.

in DV's binary discrimination tests,[4] the transits themselves are marked as gapped, and the remaining data in the flux time series are sent back to TPS to be searched for additional planet candidates.

3. ITERATIVE WHITENING AND FITTING

Once the pre-fitting processes outlined in Section 2 are carried out, the flux time series is fitted using a robust Levenberg-Marquardt fitter,[5,6] which is a modified version of the MATLAB function *nlinfit*.[7] In order to do this, the contribution of slow stellar variability must be removed, since it can be quite large compared to the transit features: for example, transits of the Earth across the sun result in a peak flux reduction (transit depth) of about 100 parts per million; solar variability over a 1-day period is approximately 10 parts per million, but over a 1-year period is closer to 1000 parts per million. A *wavelet-based whitening filter* is used to remove the variations in the flux time series which occur over timescales which are long compared to a transit. A consequence of this process is that the shape of the transit is distorted by the filter; it is therefore necessary to apply the same filter to the model flux time series used in the fit, such that the model transits and the data are properly matched to one another in shape.

The wavelet-based whitening filter is determined individually for each fit, and optimized to the noise spectrum of the given target star. Given that the purpose of the whitening filter is to remove the slow variations in the star's flux time series, it is beneficial to generate the whitening filter using a time series which contains only the slow variations and has been stripped of the transits which are to be fitted. This is accomplished by first subtracting the current best-estimate transit light curve from the flux time series, and whitening the residual; a whitened version of the transit light curve is then added back to the whitened residual flux time series, and this whitened total flux time series is used as the data which constrains the fit. In this case the whitening process depends on the current estimated transit model, but the fitted transit model depends on the whitening process. As a result, it is necessary to iteratively perform the whitening and Levenberg-Marquardt fitting until a self-consistent combination of whitener and transit model is obtained.

3.1 Initial Estimate of Transit Parameters

As described in Section 2, the initial model of the transit is derived from the TCE, which in turn is furnished by the TPS software module. The TCE contains the epoch and period of the maximum multiple event statistic, as well as the values of the multiple event statistic (MES) and the maximum single event statistic (SES). It also returns the transit duration which was used in the search: the current configuration uses 3-hour, 6-hour, and 12-hour transits to produce TCEs and returns the TCE which has the largest MES. Prior to the start of fitting, the parameters in the TCE must be converted to physical parameters for a transiting-planet solar system; these parameters are then used to seed the fit.

The conversion is simplified by assuming a circular orbit and a central transit, and by using the star radius parameter for the target star which is given in the *Kepler Input Catalog* (KIC).[8] The remaining parameters are the epoch, semi-major axis, and planet radius. The epoch is obtained directly from the TCE, and the semi-major axis can be obtained from the orbital period T, the star radius R, and the surface gravity g of the star from Kepler's Third Law:

$$a = \left(\frac{T^2 g R^2}{4\pi^2} \right)^{1/3}. \tag{1}$$

The surface gravity of the star, like the radius, is available from the KIC.

The ratio of the planet radius to the star radius is given by the product of the square root of the transit depth and a correction for limb-darkening, and therefore the planet radius can in principle be deduced from the transit depth, limb darkening parameters, and the star radius. Unfortunately the transit depth is not returned as part of the TCE, but the single event statistic can be converted to an estimate of the transit depth:

$$D \approx \mathrm{SES} \cdot \sigma_1 \cdot \sqrt{N_{\mathrm{meas}}}, \tag{2}$$

where σ_1 is the typical relative noise in a single 30 minute measurement of the flux of the target star (typically in PPM) and N_{meas} is the number of 30-minute measurements in the transit duration used in the search (*i.e.*, $N_{\mathrm{meas}} = 12$ for a 6 hour transit duration).

The depth estimate which is thus obtained from the SES is only approximately correct, but it is sufficiently accurate to use as the starting point for the fit. A minor improvement in accuracy is obtained by requiring the depth to be the minimum of D, as estimated above, and the full range of variation of the flux time series.

3.2 Whitening Filter

The top panel of Fig. 3 shows a flux time series containing a series of transits and stellar variability. The second panel shows a model series of transits which matches the actual transit depths in the flux time series. It is clear that, in order to make the functional form of these two curves match well enough to use the former as constraints for fitting the latter, it will be necessary to remove the slow component of the stellar variability while leaving the transit-like features intact. Furthermore, given that stellar variability is a non-stationary process, the frequency content of the stellar variability is itself varying with time, and any attempt to filter the flux time series must take this into account.

Given the nature of the problem, which is removal of a non-stationary non-white noise component from a time series, a joint time-frequency representation of the noise, and of the filter, is indicated; for this reason, a wavelet-based whitening filter is used for removing the stellar variability of target stars.[9] In order to prevent the whitener from removing or degrading the transit-like features of the flux time series, the following procedure is followed:

- A *residual flux time series* is formed by subtraction of the model transit from the flux time series
- The whitening filter for the fit is constructed from the residual flux time series
- The whitening filter is applied separately to the residual flux time series and the model transit, and the sum of the whitened residual flux time series and the whitened transit model is used to constrain the fit.

The whitening filter which is generated from the residual flux time series is then used in the fit, as described in the next section. In Fig. 3, the fourth panel shows the results of applying the whitening filter to the residual flux time series: as expected, the residual flux time series is converted to white noise with unit variance. The fifth panel of Fig. 3 shows the results when the whitening filter is applied to the model transit: note that the shape of the transit is distorted, such that the original transit (second panel), which was purely negative, now has positive "wings" on either side of the transit. The bottom panel of Figure 3 shows the sum of the whitened residual flux and the whitened transit model. Note that the whitener also performs a scale transformation: in the whitened domain, the dimensions are multiples of the RMS of the white noise.

3.3 Levenberg-Marquardt Fit of Transit Model

The Levenberg-Marquardt fit of the transit model uses a modified version of the MATLAB *nlinfit* function. The *nlinfit* function supports robust fitting, in which an initial fit is performed and the data points are reweighted based on the magnitude of their residuals to the fit; the fit is repeated, new weights are applied to the data, and this process of reweighting and refitting is iterated until the change in fit parameters from one iteration to the next falls below a predetermined threshold. The key modification to *nlinfit* for use in the DV fitter is a less strict convergence criterion: while the unmodified *nlinfit* continues to iterate until all parameters are stable to within 1.5×10^{-8} of their values, the modified *nlinfit* continues to iterate until all parameters are stable to within 0.5 of their estimated uncertainties. The nominal converge criterion requires a much larger number of iterations than the modified convergence criterion, and represents a case of severely diminishing returns given that the changes in parameters for the nominal convergence criterion are usually extremely small compared to the estimated parameter uncertainties.

The fit is constrained by the whitened flux time series described in the previous section. For each new set of fit parameters, an unwhitened transit model is generated and then passed through the current whitening filter; the difference between the whitened transit model and the whitened flux time series is the quantity which is minimized by the fit. The initial fit is weighted by the estimated uncertainties in the flux time series, rescaled according to the scale factor between the unwhitened and whitened domains. In the robust fitting stages, weights are a product of the rescaled uncertainties and the robust weights.

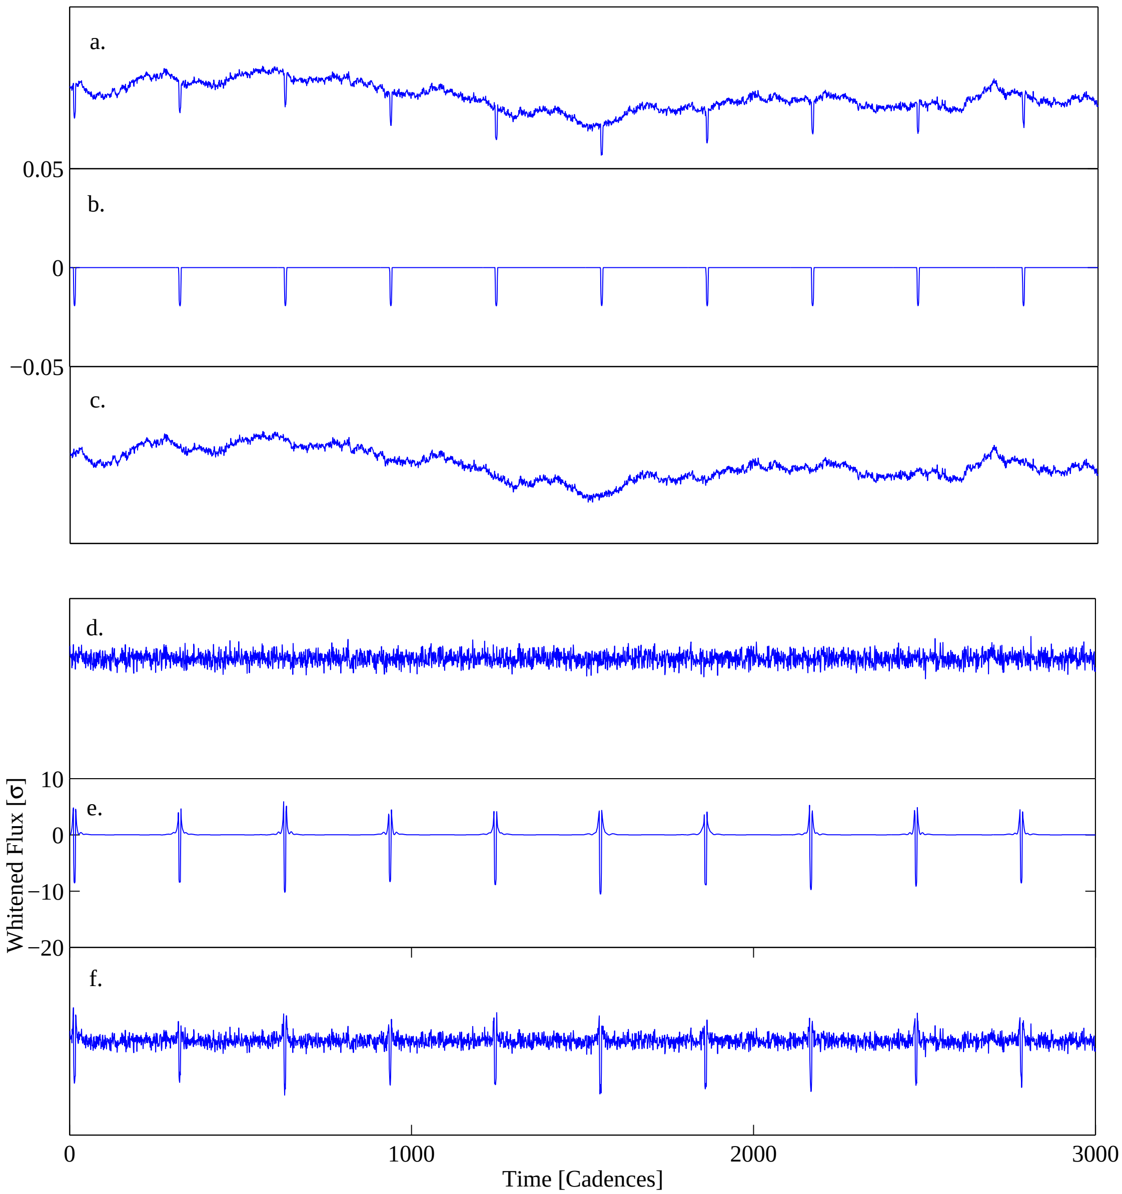

Figure 3. Whitening process for a flux time series: (a) Time series with transits and stellar variation; (b) model transit flux time series; (c) residual flux time series; (d) whitened residual flux time series; (e) whitened transit model; (f) whitened flux time series, sum of (d) and (e). Vertical scale for (a-c) dimensionless relative flux, (d-f) is multiples of the standard deviation of (d).

3.3.1 Parameterization of the transit model

For a transiting planet, there are seven physical parameters required to describe the system: transit epoch, star radius, planet radius, semi-major axis, impact parameter, eccentricity, and longitude of periastron. The *Kepler* data lacks the time resolution required to directly determine the eccentricity, and thus all orbits are modeled as circular. This leaves five free physical parameters which can be freely converted to equivalent observable parameters: epoch, orbital period, transit depth, transit duration, and transit ingress time. The ensemble of physical parameters has the advantage that any combination of valid physical parameters is itself a valid description of a transiting planet system; by contrast, there are combinations of observable parameters which are not valid as an ensemble (a trivial example of this is a system in which the transit duration exceeds the orbital period). The physical parameters are therefore the optimal ones to use for Levenberg-Marquardt fitting.

During testing and development of the DV fitter, two degeneracies in the parameters were uncovered. The first degeneracy is between the semi-major axis and the star radius: to lowest order, changing either the semi-major axis or the star radius primarily changes the observed orbital period of the system, so the fit is unable to distinguish between the two parameters. This was solved by replacing the star radius fit parameter with an orbital period fit parameter; for each new combination of parameters requested by the Levenberg-Marquardt process, the model function would use Kepler's Third Law to compute the implied star radius from the period, the semi-major axis, and the catalog value of the surface gravity, via a rearrangement of Equation 1; in this way, the transit model still uses a purely physical set of parameters even though the Levenberg-Marquardt process uses one observable parameter.

The second degeneracy is related to the impact parameter. Consider a central transit with a given period T, depth D, and duration t: one can formulate a transit with non-zero impact parameter b which has the same values of T, D, and t, but which is distinguishable from the central transit by the shape of its ingress and egress regions (essentially, the non-central transit has a longer ingress and egress time). *Kepler's* time resolution is 30 minutes and typical ingress/egress times for transiting planets in the HZ are under an hour, thus the ingress and egress times are only a weak constraint on the fit. The impact parameter is therefore only weakly constrained, and since the impact parameter is strongly covariant with the other physical parameters, inclusion of the impact parameter in the fit causes poor performance overall. To combat this, the fit is initially performed with the impact parameter held constant at zero. Once the fit has converged, the star radius in the fitted model is compared to the star radius in the KIC: a radius which is smaller than the KIC radius indicates that the data can be fitted with a model which holds the star radius constant at the KIC value and fits the impact parameter. In such cases the fit is repeated with the epoch, planet radius, semi-major axis and impact parameter used as fit parameters and the star radius held constant at the KIC value.

An additional subtlety to the parameterization is that the impact parameter is constrained to lie in the range $[0, 1]$, but the Levenberg-Marquardt algorithm implicitly requires all fit parameters to be valid over all real values. To address this mismatch, a nonlinear transformation is performed between the "internal" parameter used by Levenberg-Marquardt and the "external parameter" used in the transit model;[10] this transformation maps the range $[-\infty, \infty]$ used by Levenberg-Marquardt to the range $[-1, 1]$ in the transit model; the transit model then treats negative impact parameters as identical to their absoulte values.

The units of the fitted parameters are as follows: transit epoch, in barycentric-corrected Modified Julian Date (MJD); planet radius, in Earth radii; semi-major axis, in Astronomical Units (AU); period, in days; impact parameter is dimensionless.

3.3.2 Parameter step sizes in Levenberg-Marquardt fitter

The *nlinfit* function performs a finite-difference calculation on each of the fit parameters to determine its Jacobian. For the purposes of fitting the light curve of a transiting planet system, the main constraint on the finite-difference calculation is that the step size should be small enough that the model transits do not move by a significant fraction of their duration. This is important because the transits occupy only a small fraction of the total light curve: therefore, if the Jabobian calculation is allowed to "jump" a model transit by an interval which is comparable to the transit duration, the model can easily get into a state in which the transits in the model line up with the inter-transit intervals in the data; at this point, small changes in the transit timing have no impact on the goodness of fit, and the fitter becomes irretrievably lost.

The *nlinfit* function uses a default minimum step size of 5.05×10^{-6} of each parameter to compute its Jacobian matrix. For most of the parameters this is acceptable, but for epoch it is not. The typical epoch MJD values are around 55,000, so *nlinfit* will change the epoch by 0.33 days when computing the Jacobian. This is addressed by forcing *nlinfit* to use a minimum step size for the epoch which is 0.1 times the data sample duration, or about 3 minutes for standard *Kepler* "long cadence" data.

3.4 Convergence Criteria

As described above, the DV fitter must iterate the process of deriving a whitening filter and performing Levenberg-Marquardt fitting in the whitened domain; the iteration is necessary because the fit results depend upon the whitening filter, but the design of the whitening filter depends upon the subtraction of the transit signature from the flux time series (i.e., the whitener depends upon the fit results). Iterations of whitening filter design and model fitting cease when one of the following conditions occurs:

- The number of iterations of whitening filter design and model fitting reaches a user-selected limit (currently set to 100)
- The total time spent performing fits on the current target star reaches a user-selected limit (currently set to 9 hours)
- The change in parameter values from the previous iteration to the current one is smaller than some user-selected fraction of the estimated parameter uncertainty (currently set to 0.01).

In the event that the user has requested robust fitting, the iterative whitening-fitting process is first run to convergence without the application of robust weights. Once the non-robust fit has converged, the fitter begins a new series of whitening-fitting iterations which includes robust weights. Note that this fitting process can be extremely time-consumimg: the fitter internally iterates the Levenberg-Marquardt fit with varying weights, and the whitening-fitting loop iterates the robust fit until full internal consistency is reached. For this reason, the robust fitting process is not performed until a non-robust version of the fit has converged.

Finally, in the case in which the fitted star radius is smaller than the KIC value of the star radius, an additional set of iterations is performed in which the impact parameter is fitted and the star radius is held fixed at the KIC value. Again, the fit is allowed to converge in a non-robust manner, after which the robust fit is performed. Once the fit is complete, the final fit with fixed impact parameter is compared to the final fit with fitted impact parameter, and the fit with the lowest reduced χ^2 is returned as the best fit to the data.

For each ensemble of whitening-fitting processes (non-robust with fixed impact parameter, robust with fixed impact parameter, non-robust with fitted impact parameter, robust with fitted impact parameter), the fitter is permitted to execute up to 100 iterations of whitening and fitting. In an extreme case, the total number of iterations could reach 400. The amount of clock time allowed per target star in the fitter remains fixed at 9 hours, regardless of whether robust fitting is selected or fitting with variable impact parameter is required.

Fig. 4 shows the results of a transit model fit. The upper plot shows the whitened flux time series, the lower plot shows the same data folded at the fitted period and averaged into 30-minute wide bins.

3.5 Fitting of Odd-Numbered and Even-Numbered Transits

Once the main fit has completed, the DV fitter performs separate fits of the odd-numbered and even-numbered transits in the flux time series. This is done to provide information for the eclipsing binary discrimination tests, which are described briefly below and in detail elsewhere.[4] The all-transits fit is used to seed the odd- and even-transits fits, and the undesired transits (even-numbered transits in the odd-transits fit and vice-versa) are removed by marking their entries in the flux time series as missing.

The odd- and even-transits fits generally proceed in the same manner as the all-transits fits, with two exceptions. First, the fit parameterization of the all-transits fit (fixed or fitted impact parameter) is used for the odd- and even-transits fit. Second, depending on the total number of transits in the light curve, there may be only one transit in either the even-transits fit or in both the odd- and even-transits fits. If this is the case, then

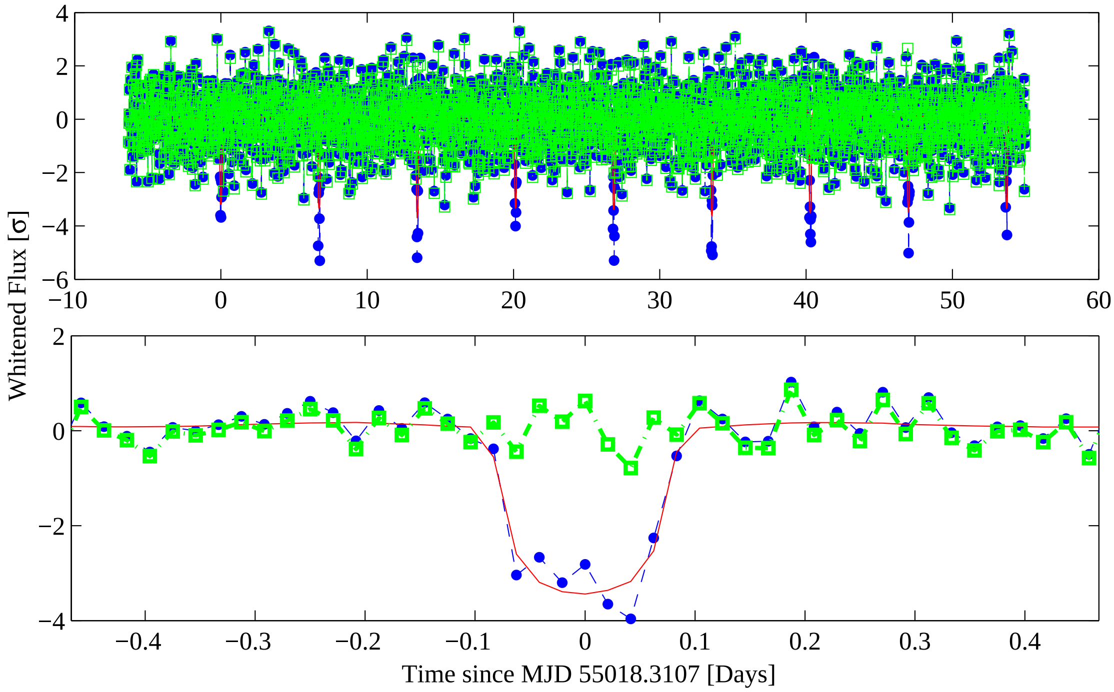

Figure 4. Sample results of the DV fitter. Top plot: whitened flux time series over the full duration of the data set. Bottom plot: whitened flux folded at the fitted period, binned and averaged to 30 minute intervals. In both plots, the original data (blue circles, dashed line), fitted light curve (red solid line), and residual time series (green squares, dot-dashed line ≡ data minus fit) are shown.

the number of parameters in the fit must be reduced by one, since the orbital period is no longer available as a fit constraint; this is accomplished by fitting only the epoch, planet radius, and semi-major axis, while holding the orbital period and the impact parameter fixed at their values determined by the all-transits fit.

4. MULTIPLE PLANET SEARCH

The TPS algorithm can only identify one planet candidate at a time, and that planet candidate will always be the one which has the largest MES; this is generally the candidate with the deepest transits. In order to locate additional planet candidates in the flux time series of a given target star, it is first necessary to remove the transit signatures of the earlier, larger candidates. This is accomplished by using the fitted transit model from the previous planet candidate to identify the timestamps which correspond to transits in the flux time series; these timestamps are then marked as missing, and the gapped flux time series is sent back to TPS to search for additional planet candidates. If additional TCEs are found which are above the detection threshold, the gapped flux time series and the new TCE are sent back into the DV fitter.

For planet candidates with extremely deep transits, it is possible for the actual transits in the data to be of longer duration than the model transits in the fit. This is especially true if the quality of the fit is not extremely good, for example when a planet model is fitted to an eclipsing binary. In such cases the transit model's identification of which data points are in transit is of limited accuracy, and the data points which are on the outskirts of a model transit can be sufficiently darkened as to trigger a transiting planet detection. This undesirable outcome is prevented by marking a number of data points adjacent to each model transit as missing data. At this time, the fitter is configured to remove data points over a time span of three transit times at each transit location; in other words, for model transits with a 10-hour duration, the multiple planet search removes the 10 hours of each model transit, and also the preceding and following 10 hours of data.

There are two additional protections against repeated detections of the same transit signature in a light curve. First, the maximum number of planet fits for each target star is currently limited to four; once four planet candidates have been fitted, DV will proceed to the next target even if the multiple planet search detects a fifth planet candidate. Second, the maximum clock time which may be used in processing any given target star is limited to 9 hours.

5. APPLICATIONS OF THE FITTED PLANET MODELS

The transit model fits from the DV fitter are used as inputs to a number of additional tests which are performed in DV. These tests are described in greater detail elsewhere,[4] and are only summarized here.

5.1 Centroid Motion Test

If the target star is actually a blend, then the photocenter will move during transits; this indicates that the transits might actually be caused by a background eclipsing binary. The centroid motion test uses the fitted transit model to determine the data timestamps which correspond to the maximum reduction in flux, and searches for a change in the photocenter location which occurs at these times. The statistical significance of the photocenter motion is then assessed.

5.2 Eclipsing Binary Tests

An eclipsing binary or background eclipsing binary can imperfectly mimic a transiting planet signature. The DV module performs a number of tests which can be used to discriminate between a planet and an eclipsing binary:

- For an eclipsing binary star with a circular orbit, the primary and secondary eclipses will be flagged by TPS as a single planet candidate. In this case, the odd-transits fit and even-transits fit will converge to different values of the transit depth. The *depth test* uses the fitted depths to determine the probability that the planet candidate is a circularized eclipsing binary.
- For an eclipsing binary star with a near-circular orbit, the interval from primary to secondary eclipse will be slightly different than the interval from secondary to primary. In this case, the fitted transit epoch of the odd-transit fit will not agree with the fitted epoch of the even-transit fit. The *epoch test* uses the fitted epochs to assess the probability that the planet candidate is actually a nearly-circularized eclipsing binary.
- For an eclipsing binary with an elliptical orbit, the primary and secondary eclipses will be identified as two distinct planets, but their fitted orbital periods will be identical. The *period test* compares the fitted periods of the all-transits fits of the planet candidates on a given target star to determine the probability that two planet candidates are actually the two eclipses of an eclipsing binary. This test also uses the estimated periods of TCEs which have been rejected from fitting due to their depth, as described in Section 2.3.

5.3 Bootstrap Analysis

At the conclusion of model fitting, it is possible to identify and remove all data points which occurred in or near a model transit. The remaining data points contain only the stellar variation and instrument noise contributions to the flux time series. These data points are therefore ideal for performing an after-the-fact *bootstrap analysis* of the fitted transits, which allows a more accurate estimate of the probability that each TCE was a result of a statistical fluctuation rather than a true astrophysical signature.

6. PERFORMANCE OF THE PLANET FITTER

The DV planet fitter was validated in an exercise which used simulated data with known ground-truth parameters. The exercise included 70 targets with an assortment of single and multiple planet systems, eclipsing binaries, and background eclipsing binaries. Out of 72 simulated true planets in the ensemble, 53 were correctly identified as planets and fitted, while 19 planets were not identified (*false negatives*). Of these, nine planets were missed because the simulated planet was too small to produce a TCE above the detection threshold; the remaining 10 false negatives were due to a number of issues in the science processing pipeline. The same exercise produced

12 *false positives*, in which non-planet signatures were mistakenly identified as planets. The vast majority of the false positives were caused by eclipsing binaries, which is an expected outcome of the DV fitter. These cases can be expected to be identified by the eclipsing binary discrimination tests in DV.

The DV planet fitter was also successfully exercised against a 90-day sample of flight data. Performance was generally good, although a few percent of the TCEs in the flight data could not be fitted successfully for reasons which are still under study.

The main fitter issue exposed by both tests was the execution time of the DV fitter. In order to perform all the fits required for the 90-day flight data set, a total of 98 DV processes running in parallel required over four days of clock time. At the time of this writing, this is the most time-consuming process in the *Kepler* processing pipeline.[11] A number of worthwhile optimizations have been identified, and will be implemented in the near future.

7. CONCLUSIONS

The DV fitter is a tool which performs automated fitting of transiting planet models to flux time series for the *Kepler Mission*. It has been successfully tested against simulated and real *Kepler* flight data with generally good results. A number of areas of potential improvement have been exposed, most significantly in the realm of execution time. In the near future, we expect to integrate DV into the *Kepler* processing pipeline and to use its results to guide selection and priorization of targets for follow-up observation.

ACKNOWLEDGMENTS

The authors would like to thank Miles Cote, Forrest Girouard, Todd Klaus, Sean McCauliff, and Bill Wohler for all of their work on the DV infrastructure and testing, Jason Rowe for many enlightening discussions on parameterization and modeling of transiting planet fits, and Bill Borucki and David Koch for their leadership of the *Kepler Mission*.

Funding for the *Kepler Mission* is provided by NASA's Science Mission Directorate.

REFERENCES

[1] Borucki, W., *et. al.*, "Kepler planet-detection mission: introduction and first results," *Science* 327, 977-980 (2010).

[2] Koch, D. G. *et. al.*, "Kepler mission design, realized photometric performance, and early science," ApJ Lett. 713(2), L79-L86 (2010).

[3] Jenkins, J. M. *et. al.*, "Transiting planet search in the Kepler pipeline," Proc. SPIE 7740, in press (2010).

[4] Wu, H. *et. al.*, "Data validation in the Kepler science operations center pipeline," Proc. SPIE 7740, in press (2010).

[5] Levenberg, K., "A method for the solution of certain non-linear problems in least squares," Q Appl. Math., 2(2), 164-168 (1944).

[6] Marquardt, D., "An algorithm for least-squares estimation of nonlinear parameters," J. SIAM Appl. Math., 11(2), 431-441 (1963).

[7] The Mathworks, MATLAB r2007a (2007).

[8] Latham, D. W. *et. al.*, "The Kepler input catalog," Proc. AAS 207, 1340 (2005).

[9] Jenkins, J. M., "The impact of solar-like variability on the detectability of transiting terrestrial planets," ApJ 575(1), 493-505 (2002).

[10] James, F., [MINUIT function minimization and error analysis reference manual], CERN, Geneva, Switzerland (1994).

[11] Middour, C. *et. al.*, "*Kepler* Science Operations Center architecture," Proc. SPIE 7740, in press (2010).